arXiv:1609.06399v1 [physics.soc-ph] 21 Sep 2016

Stochastic dynamics and the predictability of big hits in online videos

José M. Miotto,[1] Holger Kantz,[1] and Eduardo G. Altmann[1, 2]

[1] *Max Planck Institute for the Physics of Complex Systems, 01187 Dresden, Germany*
[2] *School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia*

(Dated: September 22, 2016)

The competition for the attention of users is a central element of the Internet. Crucial issues are the origin and predictability of big hits, the few items that capture a big portion of the total attention. We address these issues analyzing 10 million time series of videos' views from YouTube. We find that the average gain of views is linearly proportional to the number of views a video already has, in agreement with usual rich-get-richer mechanisms and Gibrat's law, but this fails to explain the prevalence of big hits. The reason is that the fluctuations around the average views are themselves heavy tailed. Based on these empirical observations, we propose a stochastic differential equation with Lévy noise as a model of the dynamics of videos. We show how this model is substantially better in estimating the probability of an ordinary item becoming a big hit, which is considerably underestimated in the traditional proportional-growth models.

Keywords: growth models — social media — attention economy — stochastic differential equations — L evy Stable distribution

I. INTRODUCTION.

YouTube is a representative example of online platforms in which items (videos in this case) compete for the attention of users [1–3]. The popularity of videos vary by orders of magnitude, resembling the fat-tailed distributions that have been reported in other online systems [4–6], in income and wealth [7], in finance [8], and in disciplines such as ecology, earth science, and physics [9]. The origin of such fat-tailed distributions is a century-old problem that lies at the heart of complex-systems science [10–15]. At the core of the different proposed models lies the idea that the current popularity (wealth) determines the future popularity gain (income) and enhances the inequality (rich-get-richer). Indeed, such (linear) *proportional growth* is the essential ingredient of Gibrat's law (used to describe the growth of firms [11, 13] and cities [16]), the Yule-Simon model (to model species genera [10] and language [17, 18]), scientific memes [19] and the preferential attachment model of network growth [20]. Proportional growth suggests that the big hits are very predictable because they originate from early advantages that are amplified over time.

The application of growth models to describe the popularity of online items bring new opportunities and challenges. On the one hand, due to the increasing availability of datasets, it becomes possible to compare models with an unprecedent accuracy. On the other hand, the expectations we have of the models are higher. For instance, a central question is to forecast and identify the origins of the big hits [21, 22], the most successful videos which capture most of the attention and produce most of the revenue through advertisement. To address this and other questions, the characterization of the heavy-tailed distribution of aggregated activity is not enough. One has to: (i) improve the description of the dynamics of individual items; and (ii) go beyond the average growth and analyze the stochastic fluctuations [23, 24]. The importance of these factors is illustrated in Fig. 1, where we show trajectories (views vs. time) of videos with the same early success (the same number of views, 3 days after publication). We see that trajectories quickly spread and that many trajectories with a weak start become popular in time. This suggests that big hits have a low predictability (i.e., they are hard to anticipate).

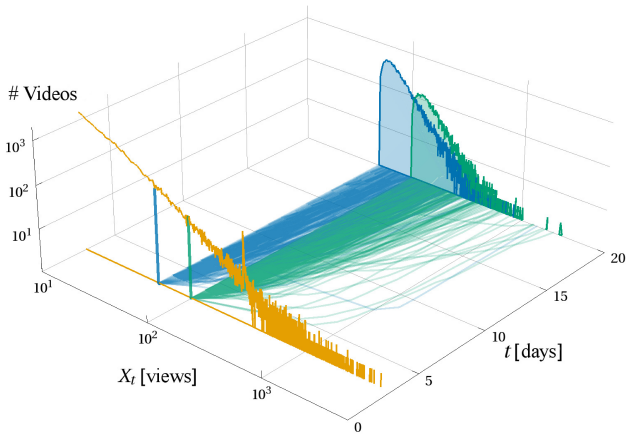

FIG. 1. Evolution of videos' views X_t as a function of the time t after publication. After $t = 3$ days the distribution of views is already heavy-tailed (orange histogram). Videos having initially the same amount of views (blue $X_{t=3} = 50$ and green $X_{t=3} = 100$) show very distinct future evolutions. On the back, histograms of the selected videos at $t = 20$ days.

In this paper, we investigate the predictability of big hits using stochastic models of individual items. Predictability is the possibility of anticipating the future based on present information and we confront the predictability expected from models to observations in the data. We compare traditional growth models to data (X_t, views over time) of more than 10 million YouTube videos. We find that previously proposed models are un-

able to correctly account for the (random) fluctuations observed in the data, which we find to be described by a Lévy-stable distribution. We propose and validate a stochastic model that explains such reduced predictability by incorporating both proportional growth and Lévy noise. This shows that, even if present, proportional growth is not the only responsible for the origin of fat-tailed distributions. Finally we show that our model substantially improves the prediction of the probability of big hits, but that unexpected big hits have an even higher probability in real data due to temporal correlations not accounted by this class of models.

II. THEORETICAL FRAMEWORK.

YouTube is a website where videos generated by third parties are shared. It is the third most visited website of all Internet. We collected more than 10 million time series of the daily number of views of videos published between Dec 2011 and Mar 2013 [25]. The number of views a video receives depends on the interplay between its content and various factors. Videos related to ongoing events are strongly influenced by their development, its media coverage, and other factors exogenous to the online activity of users. Videos are also influenced by endogenous factors, such as the sharing and recommendation in online media, generating cascades of actvity in the social network [26, 27]. Additionally, a video can be viewed by following a link from a related video, i.e. hopping through the videos' network which changes continuously according to YouTube's recommendation algorithms. The interplay and feedback between these and other factors lead to the complex dynamics we observe in the time series. Modeling specific factors [28] and differentiating between them (e.g., between exogenous and endogenous factors [26, 29, 30]) are topics of recent research. This approach is difficult to be pursued because it requires detailed information of user activities and the possibility of isolating the factors. Instead, here we aim at a coarse-grained description of the dynamics of attention in which the combination of the different factors described above are *effectively* accounted by deterministic and stochastic terms.

Let X_t be the cumulative number of views that a video received in the first t days after its release. A very general stochastic model for the growth of X_t in t is the diffusion process [23, 31, 32]

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t, \tag{1}$$

where W_t is a Wiener process ($\langle W_t \rangle = 0$ and $\langle W_t^2 \rangle = dt$), $\mu(t, X_t)$ is the average growth, and $\sigma(t, X_t)$ scales the fluctuations; an additional cutoff in dW_t is added to ensure that $dX_t > 0$. We use $dt = 1$ day, the minimum resolution of our data, so that dX_t corresponds to the number of views obtained exclusively in day $t + 1$. We consider all videos to be indistinguishable so that variations in the behavior of videos with the same X_t should be accounted by the stochastic term $\sigma(t, X_t)dW_t$. Extensions of our model could consider $\mu(t, X_t)$ and $\sigma(t, X_t)$ to depend on properties of the video and on $X_{t'}$ for $t' < t$.

III. DATA ANALYSIS.

We now analyze the data in order to identify the functions $\mu(t, X_t)$ and $\sigma(t, X_t)$. We first focus on the deterministic term of Eq. (1), $\mu(t, X_t)$. In linear proportional growth models the average growth is proportional to the views, $\mu(t, X_t) = \mu_t X_t$, where the temporal dependence on μ_t accounts for the decay in the attention gain [2]; this decay is very strong in the first weeks, so we will focus on the days up to $t = 30$.

This condition is consistent with our data: in Fig. 2(a) we see for a fixed t, that the dependence of the conditional average $\langle dX_t|X_t \rangle$ (computed in windows of N videos) with X_t is roughly a line with slope 1, a standard method to check for proportional growth [15, 33]. We now repeat the analysis for the stochastic term $\sigma(X_t, t)$ of Eq. (1). A natural proposal for σ is $\sigma(X_t, t) = \sigma_t X_t^\beta dW_t$ [32], where the β parameter allows us to model a possible fluctuations' scaling, in the form of the Taylor's Law [34]. In particular, the $\beta = 1$ case used in Ref. [23], is equivalent to $Y_t = \ln X_t$ exhibiting constant fluctuations, and corresponds to a Geometric Brownian Motion. The simplest way to evaluate the stochastic term in this context is to repeat what was done for the mean and measure the standard deviation σ in a window of N items centered around X_t [23]. This is equivalent to the standard estimation of the drift and diffusion coefficients in a Fokker-Planck Equation [35]. Results in Fig. 2(b) confirm the roughly linear scaling in the double logarithmic scale, in agreement with $\sigma(X_t, t) \propto X_t^\beta$ with $\beta \approx 1$. However, in opposite to the case of $\mu(X_t, t)$ shown in panel (a), the data show strong fluctuations across X_t and depend on the sample size N (the larger the N the larger the measured σ).

The observations above motivate us to look at the full probability distribution $\mathbb{P}(dX_t|X_t)$ [36]. In Fig. 2(c) we see in the particular histogram $\mathbb{P}(dX_3|X_3 \approx 500)$, that the distribution has heavy tails; this explains the observation that σ grows with N, i.e. $\mathbb{P}(dX_t|X_t)$ has a diverging second moment [37]. Heavy-tailed fluctuations may still be compatible with Eq. (1) if one considers that the temporal interval used in our analysis ($dt = 1$ day) is not infinitesimal (Gaussian fluctuations are expected only if $dt \to 0$). In this case, the stochastic differential equation has to be integrated, so the fluctuations predicted from Eq. (1) can be Lognormal (for $\beta = 1$) or a distribution arising from the Constant Elasticity of Variance model (CEV, for $\beta \neq 1$) [38], as shown in the Supplemental Material (SM Sec. 1). Besides Eq. (1), classical models associated with Gibrat's law (Champernowne-Gabaix or Yule-Simon) predict $\mathbb{P}(dX_t|X_t)$ to have either short tails or lognormal distributions (see SM Sec. 2). Beyond the Lognormal and CEV distributions, which follow from

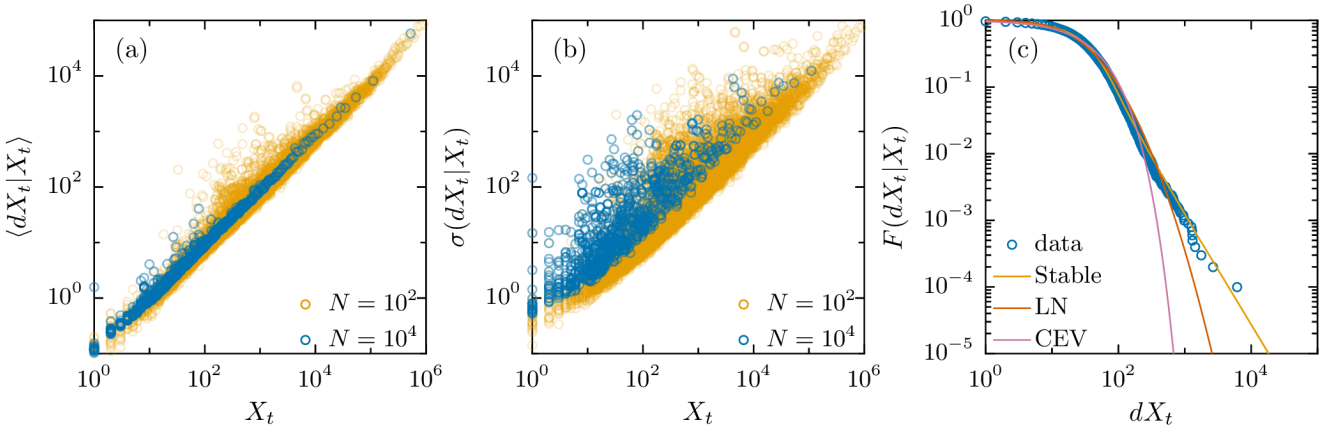

FIG. 2. Average and fluctuations in the growth of YouTube Videos. **(a)** Mean $\langle dX_t\rangle$ and **(b)** standard deviation $\sigma(dX_t)$ for videos with X_t views $t = 3$ days after publication. Both $\langle\ldots\rangle$ and σ are computed in windows centered at X_t and containing N items (see legend). **(c)** Complementary cumulative density function $F(dX_t|X_t)$ for $X_t \in [499, 513]$ and $t = 3$. Data (blue circles) is compared with fits of three distributions (S: Lévy-stable, LN: Lognormal, CEV: Constant Elasticity of Variance) and confirms the existence of heavy tails which are best described by the Lévy distribution.

Eq. (1), we consider also the Lévy-stable distribution (S) because it originates from the generalized Central Limit Theorem for variables without finite variance [39]. In Fig. 2(c) we show the fits of discretized versions of these 3 distributions to the particular histogram discussed above. The best fit is obtained by the (completely asymmetric) Lévy-stable distribution (with a difference in the Bayes Information Criterion [40], BIC, of 178 and 175, with respect to the Lognormal and CEV models). This result, which is confirmed below for different X_t and t, indicates that the fluctuations observed in the data are not compatible with the Wiener process W_t in Eq. (1), and that the analysis of the mean and standard deviation done for Fig. 2(a) and (b) may be not enough to define the functions of Eq. 1.

IV. ALTERNATIVE MODEL.

Motivated by the better fit of the Lévy distribution and by the linear scaling of μ and σ with X_t (as shown in Fig. 2), we propose as an improvement of Eq. (1) [41]

$$dX_t = \mu_t X_t dt + (a_t X_t + b_t) dL_t, \tag{2}$$

where L_t is an α-stable Lévy process, analogous to the Wiener process, except that the distribution of dL_t follows a Lévy-stable distribution with index α, asymmetry 1, location parameter 0 and scale 1 (using parameterization 1 of Ref. [42]). A cutoff in the noise term is added as above to ensure $dX_t \geq 0$, so $\langle dL_t\rangle > 0$ and $\langle dX_t\rangle$ is not given alone by the deterministic term $\mu_t X_t$ (even if dL_t is understood in the Ito sense, as we do here [43]). The parameters α, μ, a, and b depend on time t (b_t is important only for small X_t and t).

V. IMPROVED DATA ANALYSIS.

We now discuss how to determine the parameters of the two models derived from Eq. (1) and of the alternative model in Eq. (2) and to test which model best describes the data. The likelihood $\mathcal{L}_t$ of the models, for a fixed day t, is the product of the likelihoods of each distribution of dX_t conditioned on X_t with respect to the parameters of the model θ as

$$\ln \mathcal{L}_t = \sum_{X_t} \sum_{dX_t} N(dX_t, X_t) \ln f(dX_t|\theta, X_t), \tag{3}$$

where $N(dX_t, X_t)$ is the observed number of videos with a given dX_t, X_t, and f is the probability density function proposed by the models. The best parameters θ are obtained maximizing $\ln \mathcal{L}_t$ and the models are compared based on their (maximum) Likelihood, penalizing the addition of parameters (using the BIC). The distributions f we test are the same as above: Lognormal (LN) and Constant Elasticity of Variance (CEV), obtained from Eq. (1), and Lévy-stable (Stable), obtained from Eq. (2). Each of the distributions f has different parameters that depend on θ and X_t, as detailed in SM Sec. 1. This approach considers all conditional distributions $\mathbb{P}(dX_t|X_t)$, avoiding the difficulties and arbitrary choices involved in the grouping of data in windows as done in previous estimations and in Fig. 2.

Fig. 3 shows that the application of our approach to the YouTube data shows significant evidence in favor of the Lévy-stable model, Eq. (2), through the collapse of the many $\mathbb{P}(dX_t|X_t)$ in a single curve. More formally, the BIC difference of the Stable model with respect to the other models is above 10^5 for all $0 \leq t \leq 30$ (inset of Fig. 3), indicating very strong statistical support for our model (see SM Sec. 3 for details). The parameters show

a strong dependence in t in the the first week. In particular, μ_t decays in this period (reflecting a decay in the gain of views) and $\alpha_t \approx 1.75$ for $t > 5$ (see SM Sec. 4 for details on the temporal dependence of the parameters). If a model of the temporal dependence of $\alpha_t, \mu_t, a_t, b_t$ is proposed, it is possible to sum the likelihoods in Eq. (3) over t and therefore reduce the number of parameter of the models by avoiding independent fittings for each t. Altogether, these analysis support our proposal of stochastic differential equation with Lévy noise, Eq. (2), to describe the dynamics of popularity in YouTube.

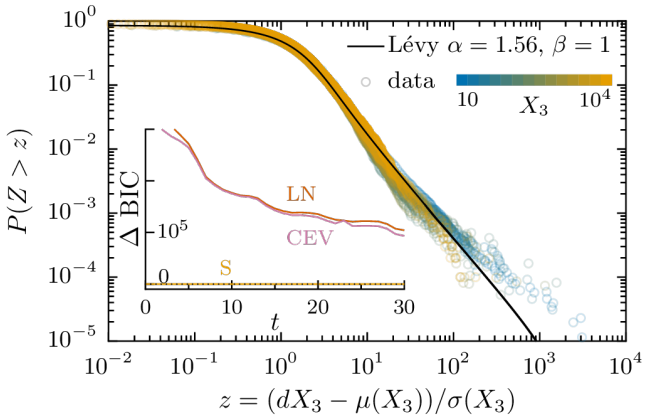

FIG. 3. Agreement of the model with respect to data. Main panel: complementary cumulative distribution of the views rescaled by the fitted parameters for $t = 3$. The histograms $\mathbb{P}(dX_3|X_3)$ are plotted as points, where each color corresponds to a different value of X_3; the black line is the Lévy-stable distribution with location 0 and scale 1. Inset: BIC difference with respect to the S model.

VI. PREDICTION OF BIG HITS.

We now focus on the estimation of the probability of an item becoming a big hit after a given time. We define as a big hit at time t the top $q\%$ videos with highest X_t ($X_t > x_t^q$). We are particularly interested in estimating the probability $P(X_t > x_t^q | X_{t_0} = x_0)$ of videos that are not big hits at time $t_0 < t$ (i.e., $x_0 < x_{t_0}^q$) becoming big hits at time t. This probability quantifies how unpredictable the system is. For instance, in a deterministic (proportional growth) model, the rank of the videos does not change and therefore such probability is zero. A positive probability is thus a measure of the deviation of such perfect predictability.

As an example, we select the videos that had 100 views one day after publication, $X_1 = 100$, which belong to a rank of $q \approx 15\%$. We are interested in the probability of these videos having $X_t \gg 100$ at $t > 1$. To obtain the expectations of the models, we computed $P(X_t|X_1 = 100)$ iteratively from $P(dX_s|X_s)$ for $s = 1, \dots t$, using $X_1 = 100$ and the t-dependent parameters estimated in the previous section. The results shown in Fig. 4(main panel) for $t = 6$ confirm that the Lévy-stable model predicts a substantially higher probability for large X_t than alternative models. In order to investigate the temporal dependence, we focus on the probability of the videos improving their rank and being by day t in the top $q = 5\%$, using the previously computed probabilities from the models and the thresholds x_t^q estimated from data. The results are summarized in Fig. 4(inset) and show that the Lévy-stable model succeeds in estimating this probability in the short-term, while for the long-term the data shows an even higher probability (mixing of ranks). The other models assign a video a substantially lower possibility of becoming a big hit, an effect of their highly predictable dynamics. The fact that our model provides a good account for short-time intervals but not in the long run suggests the existence of correlations in the attribution of views that span multiple days and that are not accounted by our assumption of an independent noise.

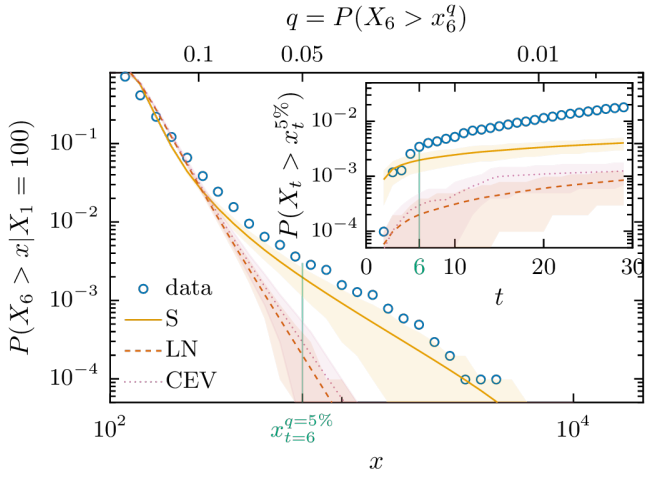

FIG. 4. Probability of videos becoming a big hit. Performance of the models evolved in time with respect to data; the selected videos had 100 views 1 day after their publication. Main panel: amount of videos that exceed a threshold x at $t = 6$. In the top axis, the quantiles q are indicated. Inset: amount of videos that enter into the 5% most viewed. Shaded areas: 95% confidence intervals, by bootstrapping.

VII. DISCUSSION AND CONCLUSION.

Our finding that the growth of views in YouTube is governed by both linear proportional growth and Lévy fluctuations has important consequences for the mathematical modeling of complex systems. First, it shows that, even if proportional growth is present, it cannot be attributed as the responsible for the origin of the heavy tails because this is a feature already present in the fluctuations. Second, the use of Gaussian-based stochastic equations, such as Eq. (1) or traditional Fokker-Planck

equations, overestimate the predictability of videos, by neglecting the mobility of popularity. We showed that better results are obtained in YouTube using a stochastic equation with Lévy noise, Eq. (2), an approach that has been previously used in Physics [41], climate research [44], and finance [8]. Our work indicates that this formalism, and possibly also kinetic equations of the fractional type [45, 46], should be considered in problems involving the dynamics of social-media items and, more generally, in models of the economy of attention.

Our results bring new insights on the attention economy of the Internet. The fact that the multiple factors affecting the popularity of videos can be effectively modeled by a Lévy-stable distribution shows that the decision of different individuals are correlated to each other and lead itself to strong fluctuations. The Lévy-stable distribution is invariant under convolution, i.e. if X_1,X_2 are stable, also $X_1 + X_2$ is stable, and therefore it is a natural attractor for the combination of multiple (fat-tailed) processes, as the ones that creates the bursty activity patterns that characterize online social media as well. One challenge for future work is to identify mechanistic models of the spreading of information on the Internet (e.g., models in which viral items spread through a social network) that are compatible with these fluctuations [23, 24, 27]. The presented analysis of fluctuations are enabled by the large availability of data in YouTube videos and we expect similar results to hold also in more general systems in which items compete for the attention of users.

ACKNOWLEDGMENTS

We thank I. Sokolov for helpful discussions.

[1] H. A. Simon, Computers, communication, and the public interest **37**, 40 (1971).
[2] F. Wu and B. A. Huberman, Proc. Natl. Acad. Sci. USA **104**, 17599 (2007).
[3] M. J. Salganik, P. S. Dodds, and D. J. Watts, Science **311**, 854 (2006).
[4] B. A. Huberman, D. M. Romero, and F. Wu, Journal of Information Science **35**, 758 (2009).
[5] J. M. Miotto and E. G. Altmann, PLoS ONE **9**, e111506 (2014).
[6] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, Sci. Rep. **2**, 335 (2012).
[7] V. Pareto, *Cours d'Économie Politique* (Librairie Droz, 1898 (1964)).
[8] B. Mandelbrot, J. Bus. **36**, 394 (1963).
[9] A. Clauset, C. R. Shalizi, and M. E. J. Newman, SIAM Rev. Soc. Ind. Appl. Math. **51**, 661 (2009).
[10] G. U. Yule, Philos. Trans. R. Soc. Lond. B Biol. Sci. **213**, 21 (1925).
[11] R. Gibrat, *Les inégalités économiques* (Recueil Sirey, 1931).
[12] B. Mandelbrot, Int. Econ. Rev. **1**, 79 (1960).
[13] H. A. Simon and C. P. Bonini, Am. Econ. Rev. **48**, 607 (1958).
[14] M. V. Simkin and V. P. Roychowdhury, Phys. Rep. **502**, 1 (2011).
[15] M. Perc, J. R. Soc. Interface **11**, 20140378 (2014).
[16] X. Gabaix, Q. J. Econ. **114**, 739 (1999).
[17] H. A. Simon, Biometrika **42**, 425 (1955).
[18] M. Gerlach and E. G. Altmann, Phys. Rev. X **3**, 021006 (2013).
[19] T. Kuhn, M. Perc, and D. Helbing, Phys. Rev. X **4**, 041036 (2014).
[20] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
[21] D. Wang, C. Song, and A.-L. Barabási, Science **342**, 127 (2013).
[22] D. J. Watts, *Everything Is Obvious: *Once You Know the Answer* (Crown Pub. Group, 2011).
[23] T. Maillart, D. Sornette, S. Spaeth, and G. von Krogh, Phys. Rev. Lett. **101**, 218701 (2008).
[24] J. Ratkiewicz, S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani, Phys. Rev. Lett. **105**, 158701 (2010).
[25] J. M. Miotto and E. G. Altmann, *Time series of social media activity. Youtube, Usenet, Stack-Overflow, PLoS ONE.*, Accessed 2015 May 21. Available at http://dx.doi.org/10.6084/m9.figshare.1160515 (2014).
[26] D. Sornette, F. Deschâtres, T. Gilbert, and Y. Ageon, Phys. Rev. Lett. **93**, 228701 (2004).
[27] D. J. Watts, Proc. Natl. Acad. Sci. USA **99**, 5766 (2002).
[28] G. Chatzopoulou, C. Sheng, and M. Faloutsos, in *INFOCOM IEEE Conference on Computer Communications Workshops, 2010* (IEEE, 2010), pp. 1–6.
[29] R. Crane and D. Sornette, Proc. Natl. Acad. Sci. USA **105**, 15649 (2008).
[30] F. Ghanbarnejad, M. Gerlach, J. M. Miotto, and E. G. Altmann, J. R. Soc. Interface **11**, 20141044 (2014).
[31] H. Risken, *Fokker-Planck Equation* (Springer, 1984).
[32] A. Mollgaard and J. Mathiesen, PLoS ONE **10**, e0123876 (2015).
[33] M. E. J. Newman, Phys. Rev. E **64**, 025102 (2001).
[34] Z. Eisler, I. Bartos, and J. Kertesz, Adv. Phys. **57**, 89 (2008).
[35] R. Friedrich and J. Peinke, Physical Review Letters **78**, 863 (1997).
[36] L. A. N. Amaral, S. V. Buldyrev, S. Havlin, M. A. Salinger, and H. E. Stanley, Phys. Rev. Lett. **80**, 1385 (1998).
[37] J.-P. Bouchaud and A. Georges, Phys. Rep. **195**, 127 (1990).
[38] R.-R. Chen and C.-F. Lee, in *Handbook of Quantitative Finance and Risk Management* (Springer US, 2010).
[39] V. M. Zolotarev, *One-dimensional stable distributions* (American Mathematical Soc., 1986).
[40] G. Schwarz, Ann. Stat. **6**, 461 (1978).
[41] A. Janicki and A. Weron, *Simulation and chaotic behavior of alpha-stable stochastic processes* (CRC Press, 1993).
[42] J. P. Nolan, *Stable distributions* (2012), ISBN 1177108605.
[43] B. Øksendal, *Stochastic differential equations* (Springer,

2003).
[44] P. D. Ditlevsen, Geophys. Res. Lett. **26**, 1441 (1999).
[45] R. Metzler, E. Barkai, and J. Klafter, Phys. Rev. Lett. **82**, 3563 (1999).
[46] D. Brockmann and I. M. Sokolov, Chem. Phys. **284**, 409 (2002).

Supplemental Material

SM-1. MODELS

We compare the data collected with the distribution predicted by a series of simple models, based on the SDEs presented in the main text. Starting from Eq. (1) of the main text, we derive two models, called Lognormal (LN) and Constant Elasticity of Variance (CEV). Instead, the Lévy-stable models (S2, S3, S4) are obtained from Eq. (2) of the main text, that differ by the amount of parameters considered.

To compare with data, we compute the distributions of X_{t+1} of the different models. Note that for the LN and CEV models, these distributions are the result of integrating Eq. (1) over a period of one day, while for the S models, this integration is not performed, i.e. we assume that in the period of one day the distribution of X_{t+1} is essentially the one of the noise.

Name	$\mathbb{P}(dX_t\|X_t)$ functional form	Parameters
LN	Lognormal	μ_t, σ_t
CEV	CEV	μ_t, σ_t, β_t
S2	Lévy-stable	α_t, μ_t, a_t
S3	Lévy-stable	α_t, μ_t a_t, b_t
S4	Lévy-stable	α_t, μ_t a_t, b_t, c_t

TABLE I. Summary of the models considered.

Lognormal (LN)

The LN model is defined by considering a linear scaling of the noise term in Eq. (1) of the main text. We have then the equation

$$dX_t = \mu_t X_t dt + \sigma_t X_t dW_t \tag{4}$$

We integrate the equation for a time equal to 1 day (where we consider μ_t and σ_t constant), such that X_{t+1} is distributed lognormally, with a probability density function

$$\mathbb{P}(X_{t+1} = x | X_t = x_0) = \frac{\exp\left(-(\log x - (\log x_0 + \mu_t - \sigma_t^2/2))^2/(2\sigma_t^2)\right)}{\sqrt{2\pi}\sigma_t x} \tag{5}$$

$dX_{t+1} = X_{t+1} - X_t$ is distributed also lognormally, since X_t is fixed, but a truncation at 0 is necessary. Since the data is distributed on the natural numbers, we discretize as well the distribution, normalizing by the sum of the PDF over its new domain.

Constant Elasticity of Variance (CEV)

If instead of Eq. SM-(4) the equation

$$dX_t = \mu_t X_t dt + \sigma_t X_t^{\beta} dW_t \tag{6}$$

is used, we have to use the distribution of the Constant Elasticity of Variance process (CEV), described in Ref. [1]. When $\beta < 1$, this distribution has the form

$$\mathbb{P}(X_{t+1} = x | X_t = x_0) = 2(1-\beta)k^{\frac{1}{2(1-\beta)}} \left(xz^{1-4\beta}\right)^{\frac{1}{4(1-\beta)}} e^{-x-z} I_{\left|\frac{1}{2(1-\beta)}\right|}\left(2\sqrt{xz}\right) \tag{7}$$

with

$$k = \frac{\mu}{\sigma^2(1-\beta)(e^{2\mu(1-\beta)}-1)}$$
$$x = k(x_0 e^{\mu})^{2(1-\beta)}$$
$$z = kx^{2(1-\beta)}$$

where I is the modified Bessel function of the first kind. The expression simplifies using the substitution $p = 2(1-\beta)$:

$$\mathbb{P}(X_{t+1} = x | X_t = x_0) = pk^{\frac{1}{p}} \left(xz^{2p-3}\right)^{\frac{1}{2p}} e^{-x-z} I_{\left|\frac{1}{p}\right|}\left(2\sqrt{xz}\right) \tag{8}$$

with

$$k = \frac{2\mu}{\sigma^2 p(e^{\mu p}-1)}$$
$$x = kx_0^p e^{\mu p}$$
$$z = kx^p$$

When $\beta > 1$, the probability function is just the same as above but multiplied by -1. Note that the β parameter is the exponent of the power law tail that the distribution has asymptotically. Here we also subtract X_t to obtain a distribution of dX_t, which we also truncate, discretize and normalize.

Lévy-stable (S)

If we use the Lévy process and linear terms for $\mu(X_t, t)$ and $\sigma(X_t, t)$, we have the SDE:

$$dX_t = (\mu_t X_t + c_t)dt + (a_t X_t + b_t)dL_t \tag{9}$$

This is the most general form, which we call S4. The model S3 has $c_t = 0$ and the model S2 has $c_t = 0$ and $b_t = 0$. dX_t is Lévy-stable distributed with location parameter $m = \mu_t x_0 + c_t$, scale parameter $s = a_t x_0 + b_t$, asymmetry $\beta_L = 1$ and its tail decays as an α power of dX_t. These parameters correspond to the parametrization 1 of Ref. [2], where the characteristic function of dX (there is no explicit form of the Lévy probability distribution function), $\phi_{dX}(k)$ is given by

$$\log \phi_{dX}(k) = \begin{cases} imk - s^\alpha |k|^\alpha \left[1 + i\beta_L \tan\left(\frac{\pi\alpha}{2}\right) \text{sign}(k)\right] & \alpha \neq 1 \\ imk - s|k| \left[1 + i\beta_L \frac{2}{\pi} \text{sign}(k) \log(|k|)\right] & \alpha = 1 \end{cases} \tag{10}$$

To get the distribution $\mathbb{P}(X_{t+1} = x | X_t = x_0)$, this characteristic function has to be transformed to the real space, translated on the X_{t+1} axis by an amount x_0, and then truncated at $X_{t+1} = x_0$, discretized and normalized, as the distributions of Eq. 2 and Eq. 5.

Numerically, the Lévy probability distribution is computed from its definition as follows:

(i) the characteristic function (its Fourier transform) is inverted numerically on a grid in$\alpha \in (0.5, 2)$ and $\beta \in (0, 1)$ (with a resolution of 0.05, values of α below 0.5 are very unlikely and for $\beta < 0$ the distribution can be computed from the one of $-\beta$ using symmetry);

(ii) for general values α, β, we compute the distribution as an interpolation of the values on the grid (using the Catmull-Rom cubic splines).

(iii) the numerical intergration often becomes unstable in the tails of the distribution (large x). In order to avoid this problem, we use the power-law approximation described in Ref. [33] of the main text to describe the distribution beyond a threshold.

We provide the code of this procedure in the package PyLevy [3]. It contains routines to compute the PDF of the Lévy distribution and to fit it.

The models considered are summarized in Tab. I

SM-2. FLUCTUATIONS EXPECTED FROM EXISTING MODELS

Here we discuss the form that $\mathbb{P}(dX_t|X_t)$ has for the classical linear proportional growth models. There are basically two schemes of implementing linear proportional growth in order to get heavy-tailed distributions:

- Scheme 1: Champernowne [4] introduces a lower positive limit to X. A master equation is defined to regulate the transitions to different states (amount of views), which eventually leads to a stationary distribution with a power-law tail. This argument was formalized and popularized in Refs. [5, 6], using a linear SDE (as Geometric Brownian Motion, GBM) which, in the limit of long time, converges to a heavy tailed distribution; note that in this scheme, all items start with the same initial condition.

- Scheme 2: Yule and Simon [7, 8] design a scheme where views are added to different items while at the same time new items are introduced, resulting in a power law distribution. This is basically the model known as *preferential attachment* [9] in the context of network growth. Here the items start in different conditions, since the system is growing in the number of items, hence the first ones are privileged.

We focus on the transition probability $\mathbb{P}(dX_t|X_t)$. Scheme 1 (GBM) is a mechanism that leads to a heavy-tailed distribution asymptotically, but also relies on the possibility of negative growth rates dX_t, which is not realistic in the context of videos' views, and more generally in the context of cumulative allocation of attention. For an infinitesimal increase of time dt, the distribution is Normal, but if a finite time interval is considered, the integration of the model results in a Lognormal distribution.

The Scheme 2, instead, is fundamentally different and can be thought of as a Polya Urn process where at a given time t_0, a number of views N is assigned to a set of videos M that has exactly x_0 views already. The probability of assigning a view to a particular video is, of course, proportional to the amount of views each video has. The distribution $\mathbb{P}(dX_t|X_t)$ of views among the videos is of the Beta Binomial type, with means

$$\mathbb{E}_{BB}(dX_t|X_t = x_0) = \frac{N}{M} \tag{11}$$

and variance

$$\mathbb{V}_{BB}(dX_t|X_t = x_0) = \frac{N(M-1)(Mx_0+N)}{M^2(Mx_0+1)} \tag{12}$$

and can be roughly approximated by a Normal distribution with same mean and variance. The variance scales in two regimes because of the $Mx_0 + N$ term: when $x_0 \lesssim N/M$, $\sigma \propto x_0^{-1/2} N/M$, and when $x_0 \gtrsim N/M$, $\sigma \propto \sqrt{N/M}$. Notice though, that the amount of views allocated, N, is not independent. In fact, since growth is linear, we expect $N \propto x_0 M$, so we have in this particular case $\sigma \propto x_0^{1/2}$.

In conclusion, for Scheme 1, we expect normal or lognormal distributions of dX_t, depending on the choice of whether integrating over a finite time or not, while for Scheme 2 we expect approximately normal distributions, with variance scaling as X_t.

SM-3. MODEL SELECTION

Parameters' estimation

We select the best parameters for each model by maximizing the likelihood of the model as a whole, as computed in Eq. (3) of the main text.

Since the distributions $\mathbb{P}(dX_t \mid X_t)$ are heavy tailed, the standard deviation is not a good measure of the fluctuations of the SDE. The solution is to fit the full distribution, and maximum likelihood estimation is the standard technique to do it. However, applying maximum likelihood estimation to each distribution we encounter other problems. The estimation is quite sensitive to the fluctuations where most data is located, and for this reason it can fail when it comes to the determination of the tail exponent if no threshold is set [10]; this causes many histograms that seem to have the same tail to be fitted with very different values of α. On another side, once the parameters of the fits for each $\mathbb{P}(dX_t \mid X_t)$ are extracted, it is not clear how to obtain the best parametrization of them, because of the presence of Lévy-distributed errors; the main limitations here are the truncation of fluctuations at $dX_t = 0$ (not to be confused with the lower positive limit on X_t of Scheme 1 described in Sec. 1 of the SI), and the discretization of the dX_t variable, which prevents a fit on the rescaled variable $z = (dX_t - \mu(X_t))/\sigma(X_t)$.

The choice of maximizing the likelihood that we propose in Eq. (3) of the main text avoids these problems in a natural way, by using maximizing the likelihood of the all the data in the same process. The maximization of the likelihood is performed by minimizing $-\log \mathcal{L}$ numerically through the Nelder-Mead algorithm, implemented in the Python package `scipy` as `optimize.fmin`. The optimization is repeated 50 times with different, random, initial conditions, in order to avoid getting stuck in a local minimum of the likelihood. Finally, the minimum of all the minima found is chosen.

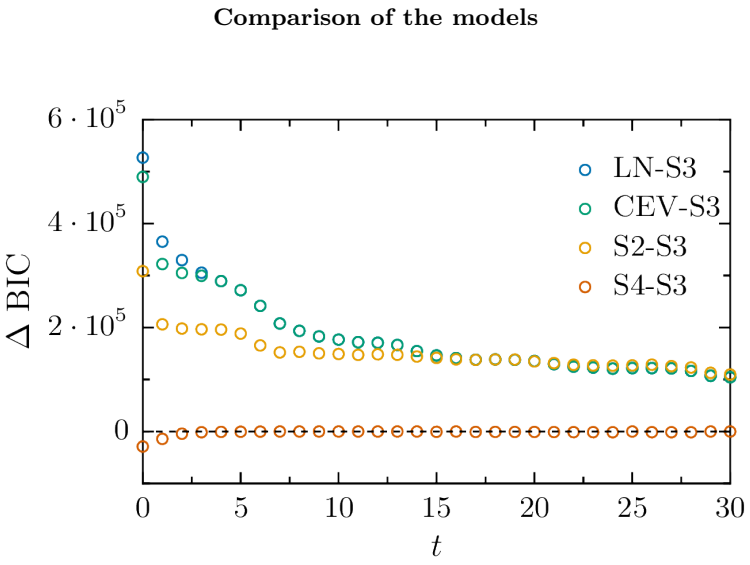

FIG. 5. Difference of the BIC of the S3 model with respect to the others, for different times.

We compare the models with the Bayes Information Criterion (BIC [11]), a way of penalizing the increase of likelihood due to the addition of parameters:

$$\mathrm{BIC} = -2 \ln \mathcal{L} + k \ln N \tag{13}$$

where $\mathcal{L}$ is the likelihood (as computed in Eq. (3) of the main text), N is the total amount of data and k is the total amount of parameters.

In Fig. SM5 we plot the difference of the BIC of model $S3$ and other model M, Δ BIC = BIC(M) − BIC$(S3)$. A higher value of the BIC difference indicates a higher support for the S3 Model.

Overall, the S3 model is the best. While the S2 model is already performing better than the LN, the difference with S3 is considerably larger. Instead, the difference of BIC between S3 and S4 is not so big, and sometimes the value of the parameters of the S4 model coincide with the S3 (c_t close to 0). So while the addition of the b_t term to the S2 model is of crucial importance, the addition of c_t is not improving the model in the same way.

SM-4. DEPENDENCE OF PARAMETERS WITH RESPECT TO TIME t

The parameters of the S3 model (Eq.(2) of the main text) are explicitly time dependent. In Fig. SM6 the values of these are shown for the first 30 days after the publication of the videos. Two main features are worth noticing. First, the parameters seem to relax to a particular value after the first week, including the exponent of the Lévy distribution, α_t. Second, the value of μ_t becomes negative; at first sight it can seem that is in contradiction with the positive slope of $\langle dX_t \mid X_t \rangle$ (see Fig. 2(a) of the main text), however, this is not the case, because the truncation of the Stable distribution at $dX_t = 0$ makes the average of the distribution **not** to be $\mu_t(X_t)$. If the average is computed from the fitted distribution, then the real average is recovered.

[1] R.-R. Chen and C.-F. Lee, in *Handbook of Quantitative Finance and Risk Management* (Springer US, 2010).

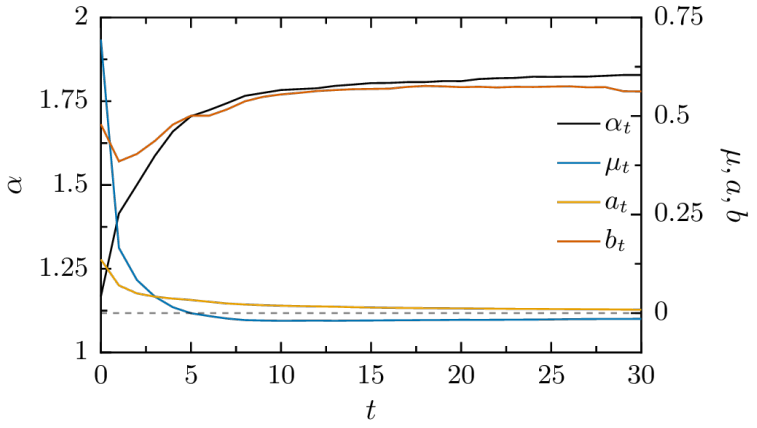

FIG. 6. Evolution of the parameters for the proposed model, Eq. (2) of main text.

[2] J. P. Nolan, *Stable distributions* (2012), ISBN 1177108605.
[3] J. M. Miotto, *pylevy*, `http://dx.doi.org/10.5281/zenodo.53787` (2016).
[4] D. G. Champernowne, Econ. J. pp. 318–351 (1953).
[5] H. Kesten, Acta Mathematica **131**, 207 (1973).
[6] X. Gabaix, Q. J. Econ. **114**, 739 (1999).
[7] G. U. Yule, Philos. Trans. R. Soc. Lond. B Biol. Sci. **213**, 21 (1925).
[8] H. A. Simon, Biometrika **42**, 425 (1955).
[9] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
[10] A. Clauset, C. R. Shalizi, and M. E. J. Newman, SIAM Rev. Soc. Ind. Appl. Math. **51**, 661 (2009).
[11] G. Schwarz, Ann. Stat. **6**, 461 (1978).